PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of:
Jerome J. CARTMELL, et al.

Appl. No.: 10/812,291

Filed: March 29, 2004

For: **MIRRORED MEMORY**

Art Unit: 2189

Examiner: VO, Thanh Duc

Atty. Docket: EMS-06601

DECLARATION UNDER 37 CFR 1.131

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

1. I understand that the originally-named inventors of the above-captioned U.S. Patent Application No. 10/812,291, filed in the United States on March 29, 2004 (hereinafter "the present patent application") are Jerome J. Cartmell, Qun Fan, Steven T. McClure, Robert DeCrescenzo, Haim Kopylovitz and Eli Shagam. I am one of the originally-named inventors.

2. I have assigned my interest in the present patent application to EMC Corporation.

3. I understand that U.S. Patent Application Pub. No. 2004/0205384 A1 to Lai et al., filed in the United States on February 18, 2004, and published October 14, 2004, has been cited by the Examiner against the present patent application.

4. Attached hereto as Exhibit A is a copy of an internal EMC document prepared by the inventors entitled "Symm7 Mirrored Memory Conceptual Review" containing presentation materials addressing hardware implementation, software requirements and limitations, and software design concepts concerning Symm7 Mirrored Memory. The dates of Exhibit A have been redacted; however, all of the redacted dates are from a time prior to February 18, 2004.

5. Specifically, pages 3-5 of the "Symm7 Mirrored Memory Conceptual Review" document of Exhibit A address "Mirrored Memory Hardware Implementation," pages 11-14 address software design concepts including "6 Different Mirrored Memory System States" and "Memory Access Functional Layering," and pages 20-25 address "Failover & Synchronization State."

6. Exhibit A, and particularly the above-noted portions, indicates that, before February 18, 2004, I and the other inventors conceived of and developed the invention set forth in the present patent application and including at least the following features:

a) "A method of accessing data memory, comprising:

writing data to a first memory location and to a second memory location in response to a request to write data to a memory address that corresponds to both locations, wherein the first and second memory locations are mirrored;

in response to a request to read data from the memory address, reading data from the first memory location or the second memory location based on load balancing; and

accessing data from the second memory location in response to a request to access data at the memory address when memory hardware corresponding to the first memory location has failed."

b) "Computer software, stored in a computer-readable medium, that accesses data memory, comprising:

executable code that writes data to a first memory location and to a second memory location in response to a request to write data to a memory address that corresponds to both locations, wherein the first and second memory locations are mirrored;

executable code that reads data from the first memory location or the second memory location based on load balancing in response to a request to read data from the memory address; and

executable code that accesses data from the second memory location in response to a request to access data at the memory address when memory hardware corresponding to the first memory location has failed."

c) "A data storage device, comprising:

a plurality of disk drives;

an internal volatile memory; and

a plurality of directors coupled to the memory, wherein some of the directors are coupled to the disk drives and some of the directors allow external access to the data storage device and wherein each of the directors access the memory by writing data to a first memory location and to a second memory location in response to a request to write data to a memory address that corresponds to both locations, wherein the first and second memory locations are mirrored, in response to a request to read data from the memory address, the directors read data from the first memory location or the second memory location based on load balancing, and the directors access data from the second memory location in response to a request to access data at the memory address when memory hardware corresponding to the first memory location has failed."

7. Further, attached hereto as Exhibit B is a copy of e-mail exchanges by and among the inventors concerning development of the subject matter of the present patent application. The dates of these documents have been redacted; however, all of the redacted dates are from a time between February 18, 2004, and the March 29, 2004, filing date of the present patent application.

8. All of the work that is reflected in the contents of Exhibits A and B was performed in the US or other WTO or NAFTA country.

9. Furthermore, during the period between February 18, 2004, and the March 29, 2004, filing date of the present patent application, I and the other inventors received communications from and responded to EMC's outside patent counsel concerning drafts of the present patent application that were being prepared by the outside patent counsel.

10. To the best of my knowledge and belief, I and the other inventors are the only individuals known to me to have conceived of and developed the above-noted features, and we were diligent from a time prior to February 18, 2004, until the invention was constructively reduced to practice by filing the present patent application on March 29, 2004.

11. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements are made with knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the above-captioned Application for Patent or any patent issuing thereon.

Respectfully submitted,

Date: 12/13/2006

Jerome J. Cartmell

Date: 12/13/2006

Qun Fan

Date: 12/13/06

Steven T. McClure

Date: 12/13/06

Robert DeCrescenzo

Date: _____

Haim Kopylovitz

Date: _____

Eli Shagam

7. Further, attached hereto as Exhibit B is a copy of e-mail exchanges by and among the inventors concerning development of the subject matter of the present patent application. The dates of these documents have been redacted; however, all of the redacted dates are from a time between February 18, 2004, and the March 29, 2004, filing date of the present patent application.

8. All of the work that is reflected in the contents of Exhibits A and B was performed in the US or other WTO or NAFTA country.

9. Furthermore, during the period between February 18, 2004, and the March 29, 2004, filing date of the present patent application, I and the other inventors received communications from and responded to EMC's outside patent counsel concerning drafts of the present patent application that were being prepared by the outside patent counsel.

10. To the best of my knowledge and belief, I and the other inventors are the only individuals known to me to have conceived of and developed the above-noted features, and we were diligent from a time prior to February 18, 2004, until the invention was constructively reduced to practice by filing the present patent application on March 29, 2004.

11. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements are made with knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the above-captioned Application for Patent or any patent issuing thereon.

Respectfully submitted,

Date: _____

Jerome J. Cartmell

Date: _____

Qun Fan

Date: _____

Steven T. McClure

Date: _____

Robert DeCrescenzo

Date: 12/13/2006

Haim Kopylovitz

Date: _____

Eli Shagam

7. Further, attached hereto as Exhibit B is a copy of e-mail exchanges by and among the inventors concerning development of the subject matter of the present patent application. The dates of these documents have been redacted; however, all of the redacted dates are from a time between February 18, 2004, and the March 29, 2004, filing date of the present patent application.

8. All of the work that is reflected in the contents of Exhibits A and B was performed in the US or other WTO or NAFTA country.

9. Furthermore, during the period between February 18, 2004, and the March 29, 2004, filing date of the present patent application, I and the other inventors received communications from and responded to EMC's outside patent counsel concerning drafts of the present patent application that were being prepared by the outside patent counsel.

10. To the best of my knowledge and belief, I and the other inventors are the only individuals known to me to have conceived of and developed the above-noted features, and we were diligent from a time prior to February 18, 2004, until the invention was constructively reduced to practice by filing the present patent application on March 29, 2004.

11. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements are made with knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the above-captioned Application for Patent or any patent issuing thereon.

Respectfully submitted,

Date: _____ _____
Jerome J. Cartmell

Date: _____ _____
Qun Fan

Date: _____ _____
Steven T. McClure

Date: _____ _____
Robert DeCrescenzo

Date: _____ _____
Haim Kopylovitz

Date: Dec 13 2006 _____
Eli Shagam

where information lives

EMC²
where information lives

Symm7 Mirrored Memory Conceptual Review

Jerry Cartmell
Bob DeCrescenzo
Chun Fan
Haim Kopylovitz
Steve McClure
Eli Shagam

EXHIBIT A

EMC²
where information lives

Presentation Overview

- Hardware Implementation
- Software Requirements & Limitations
- Software Design Concepts
- Project Status
- Q & A

EMC²
where information lives

Hardware Implementation

- Similar dual write implementation as was available (but not used) in Symm6
- Memory boards have fixed mirrored board and director port pairs. For example:
  - Memory board 0 is always paired with memory board 1
  - Director port 0 is always mirrored with director port 4

EMC²
where information lives

Mirrored Memory Hardware Implementation

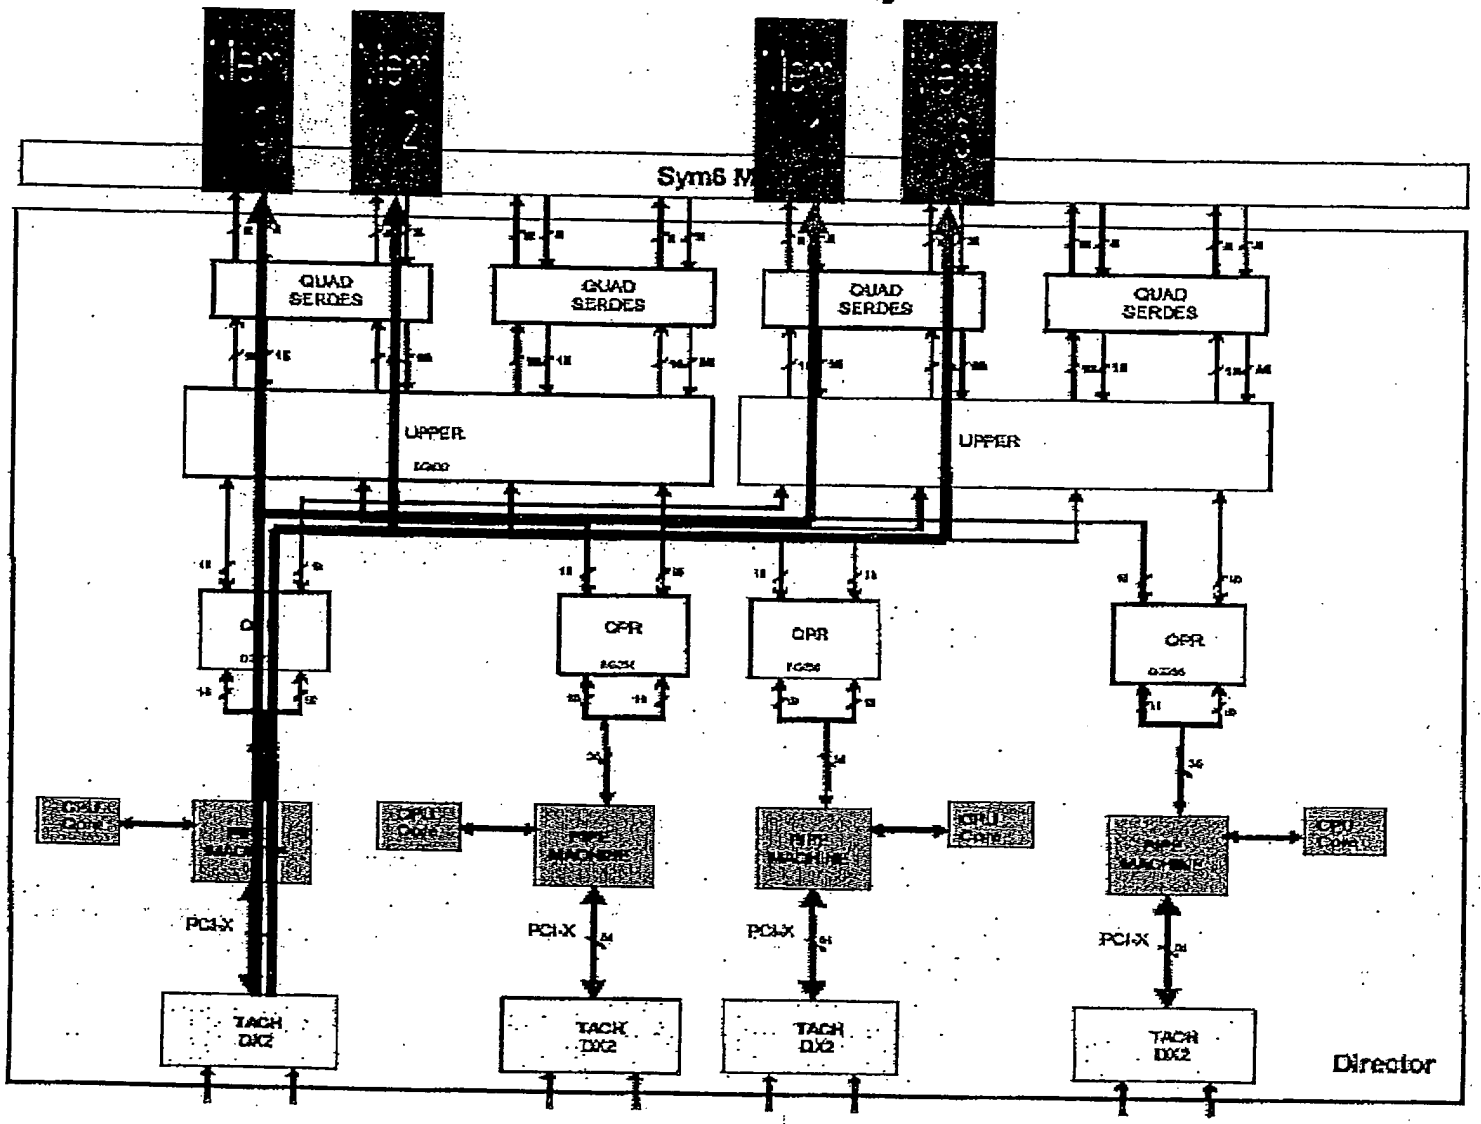

EMC²
where information lives

Hardware Implementation

- Similar dual write implementation as was available (but not used) in Symm6
- Memory boards have fixed mirrored board and director port pairs. For example:
  - Memory board 1 is always paired with memory board 2
  - Director port 0 is always mirrored with director port 4
- No global hardware switch for dual write on or off. HW Dual write is on all the time by default
- Can turn off dual write on a per I/O basis

EMC²
where information lives

Hardware Anomalies

- Hardware will not support dual writes on atomic operations (MCMS) . Software will have to do 2 separate accesses.
  - Estimated performance penalty on I/Os per sec
    - FA -- 10.6%
    - FICON -- 11.8%
    - DA -- 8.3%
- In the case of a failure in 1 transfer of the dual write, the hardware will stop both transfers.

EMC²
where information lives

1 failure kills both transfers for all I/Os In progress

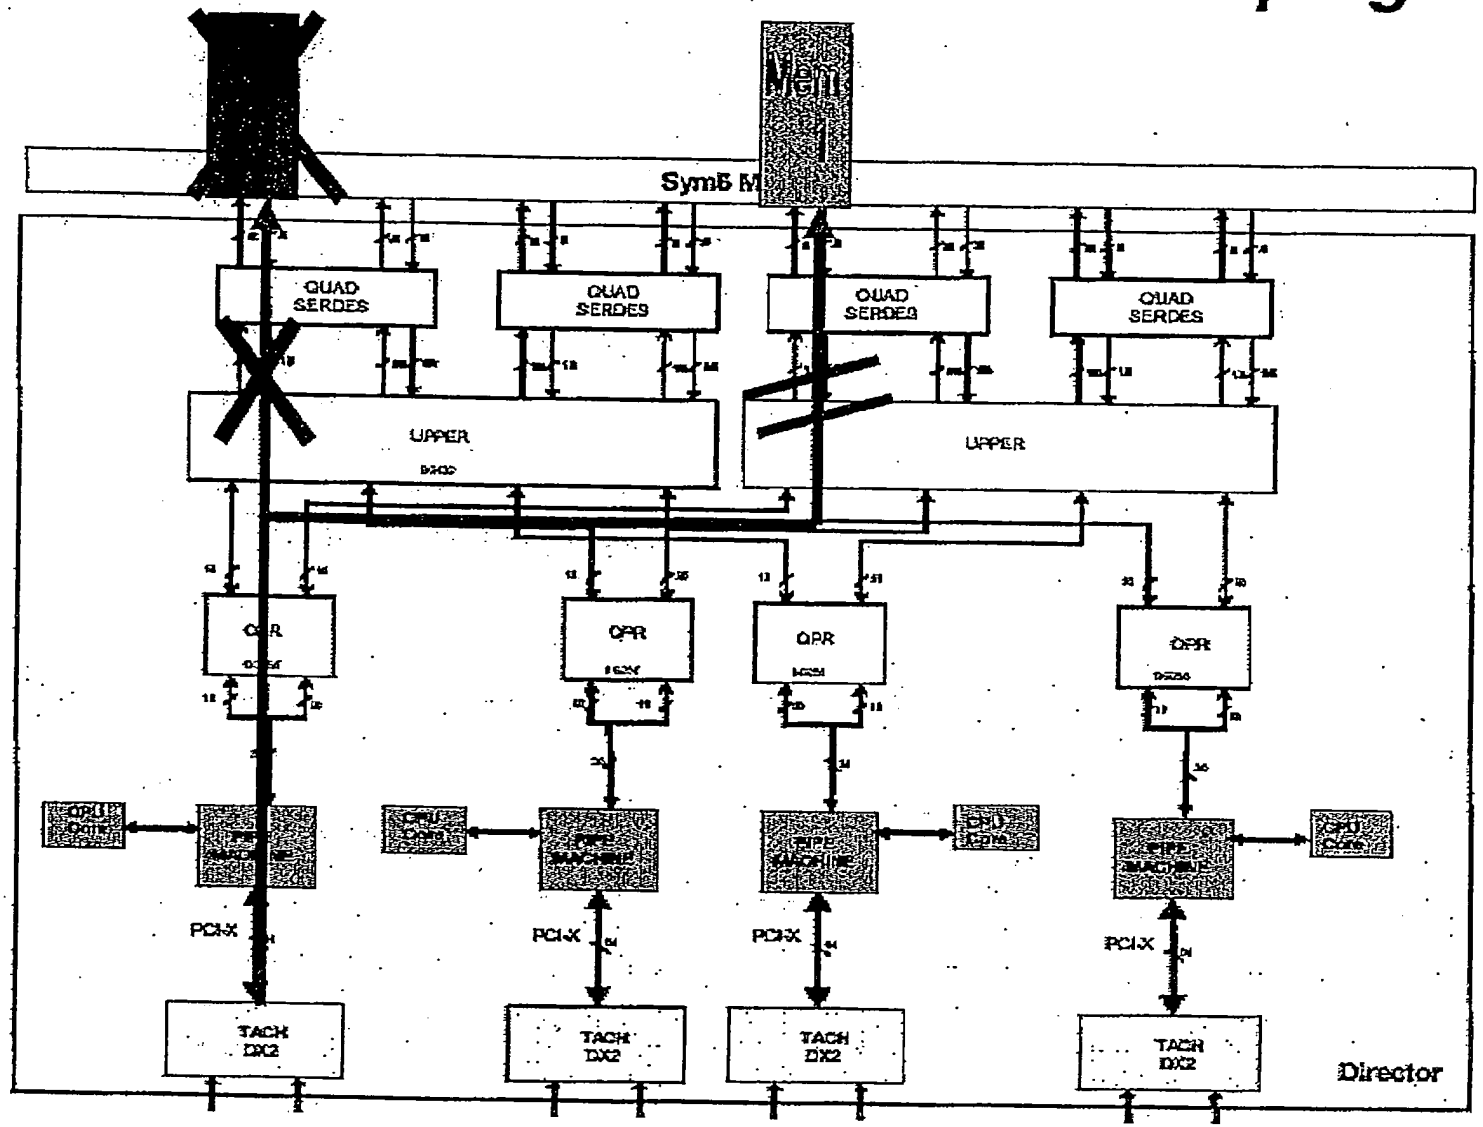

EMC²
where information lives

Software Requirements & Limitations

- Software only supports mirrored memory configurations.
- Software shall do FULL mirroring of global memory mirroring. There will not be any separate Read & Write cache.
- Software shall only failover an individual memory board. It will not failover all memory boards in the system on a single failure.
- Error failover time shall be no worse than the time it takes to do an online code upgrade.

EMC²
where information lives

Software Requirements & Limitations (Cont)

- All mirrored memory knowledge and handling is done in the HW core services layer.
- Software will no longer support things like port failover or multibit recovery. All hard memory failures will invoke mirrored memory failover.
- Memory board replacements should be done with minimum User & Symmwin intervention – "Ease of Use"
- The software will continue to write user data to an unprotected memory board.

EMC²
where information lives

Software Design Concepts

EMC²
where information lives

6 Different Mirrored Memory System States

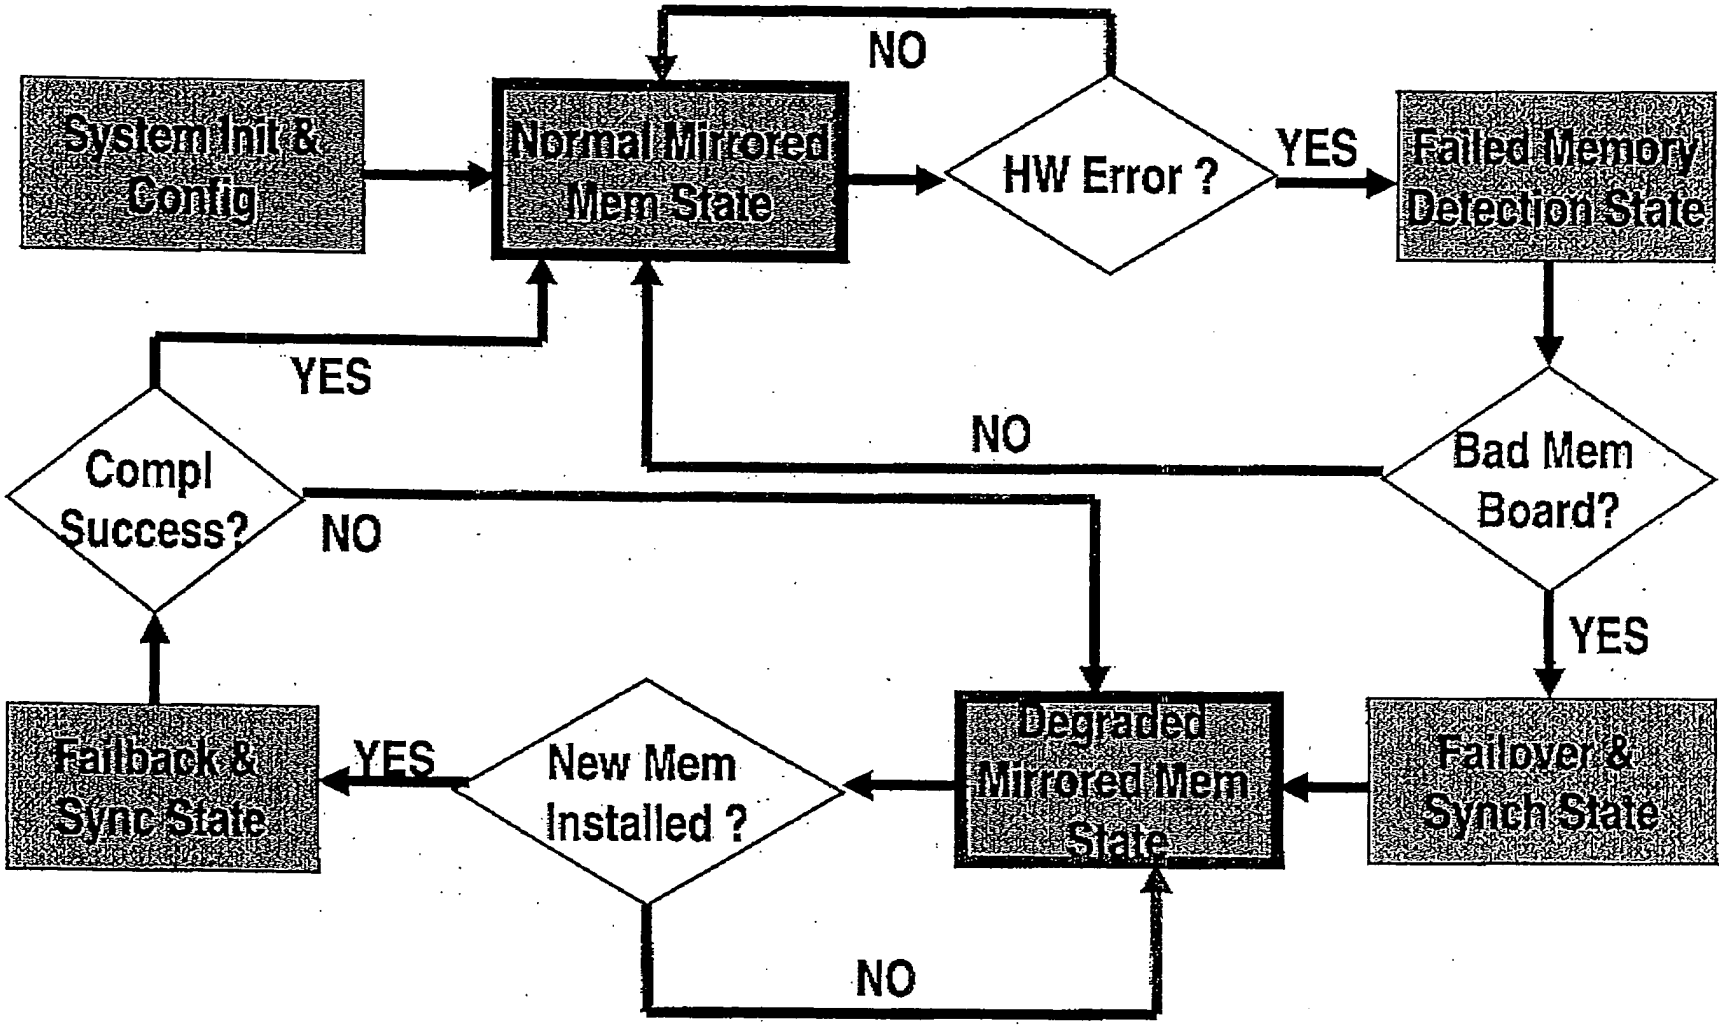

EMC²
where information lives

6 Different Mirrored Memory System States

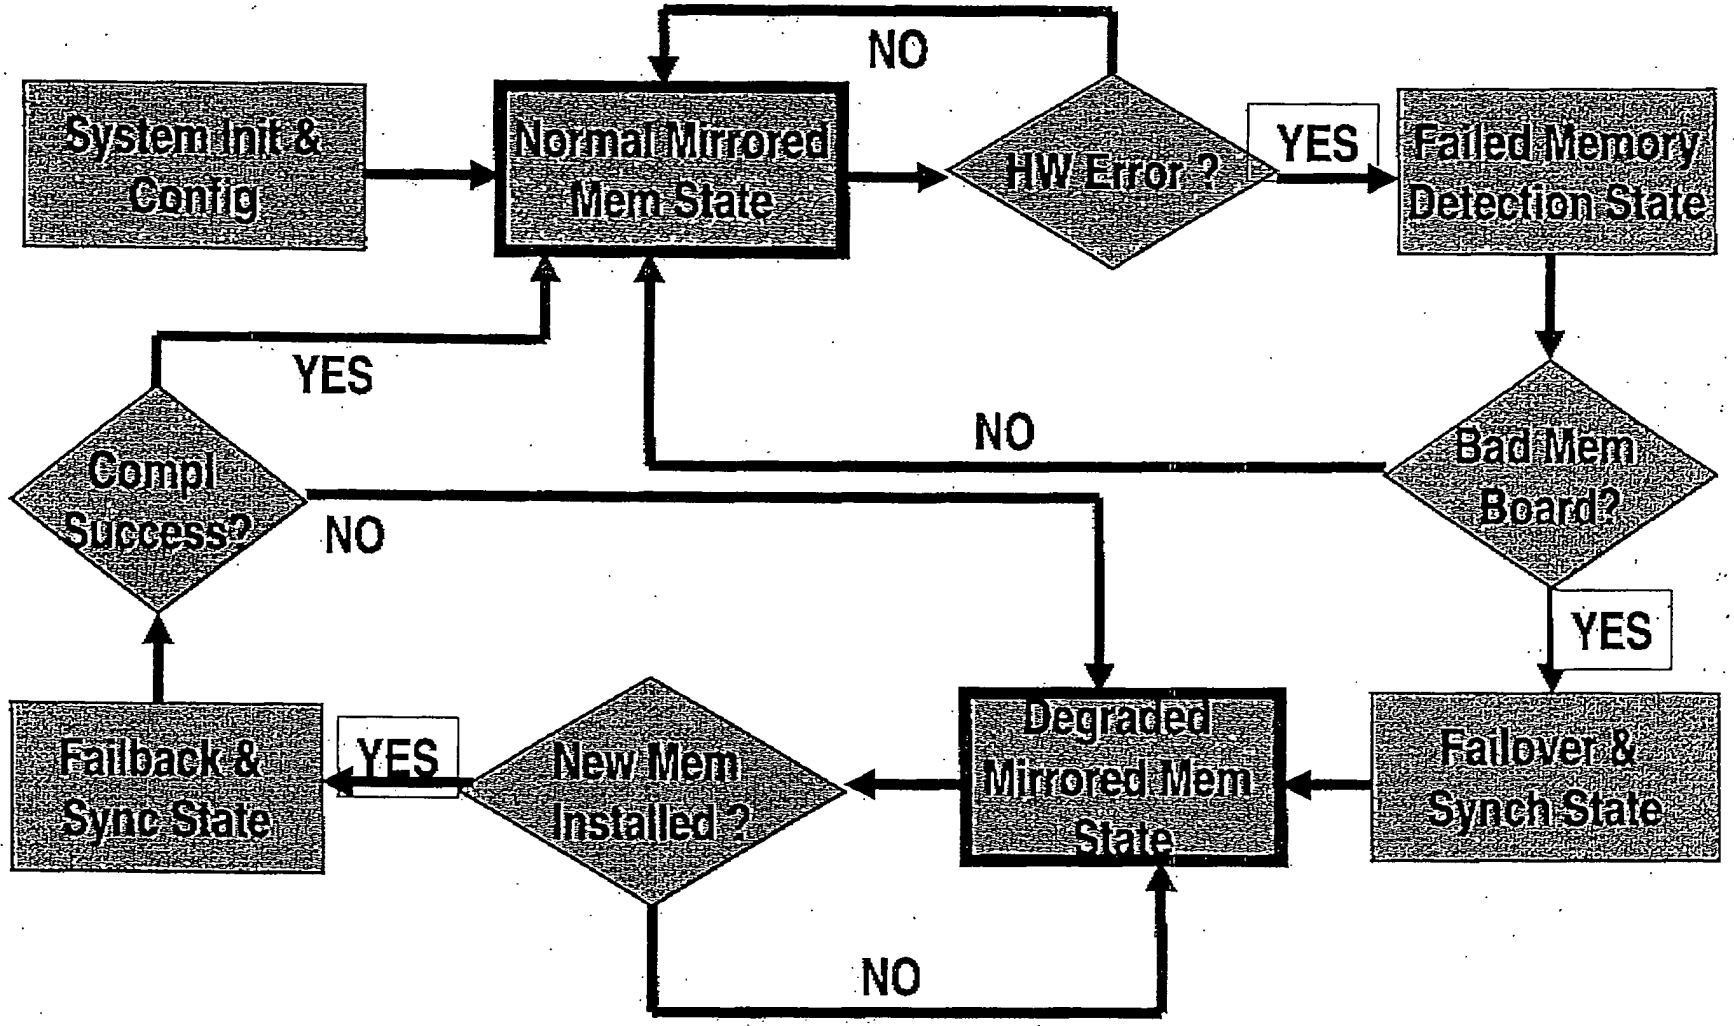

EMC²
where information lives

Memory Access Functional Layering

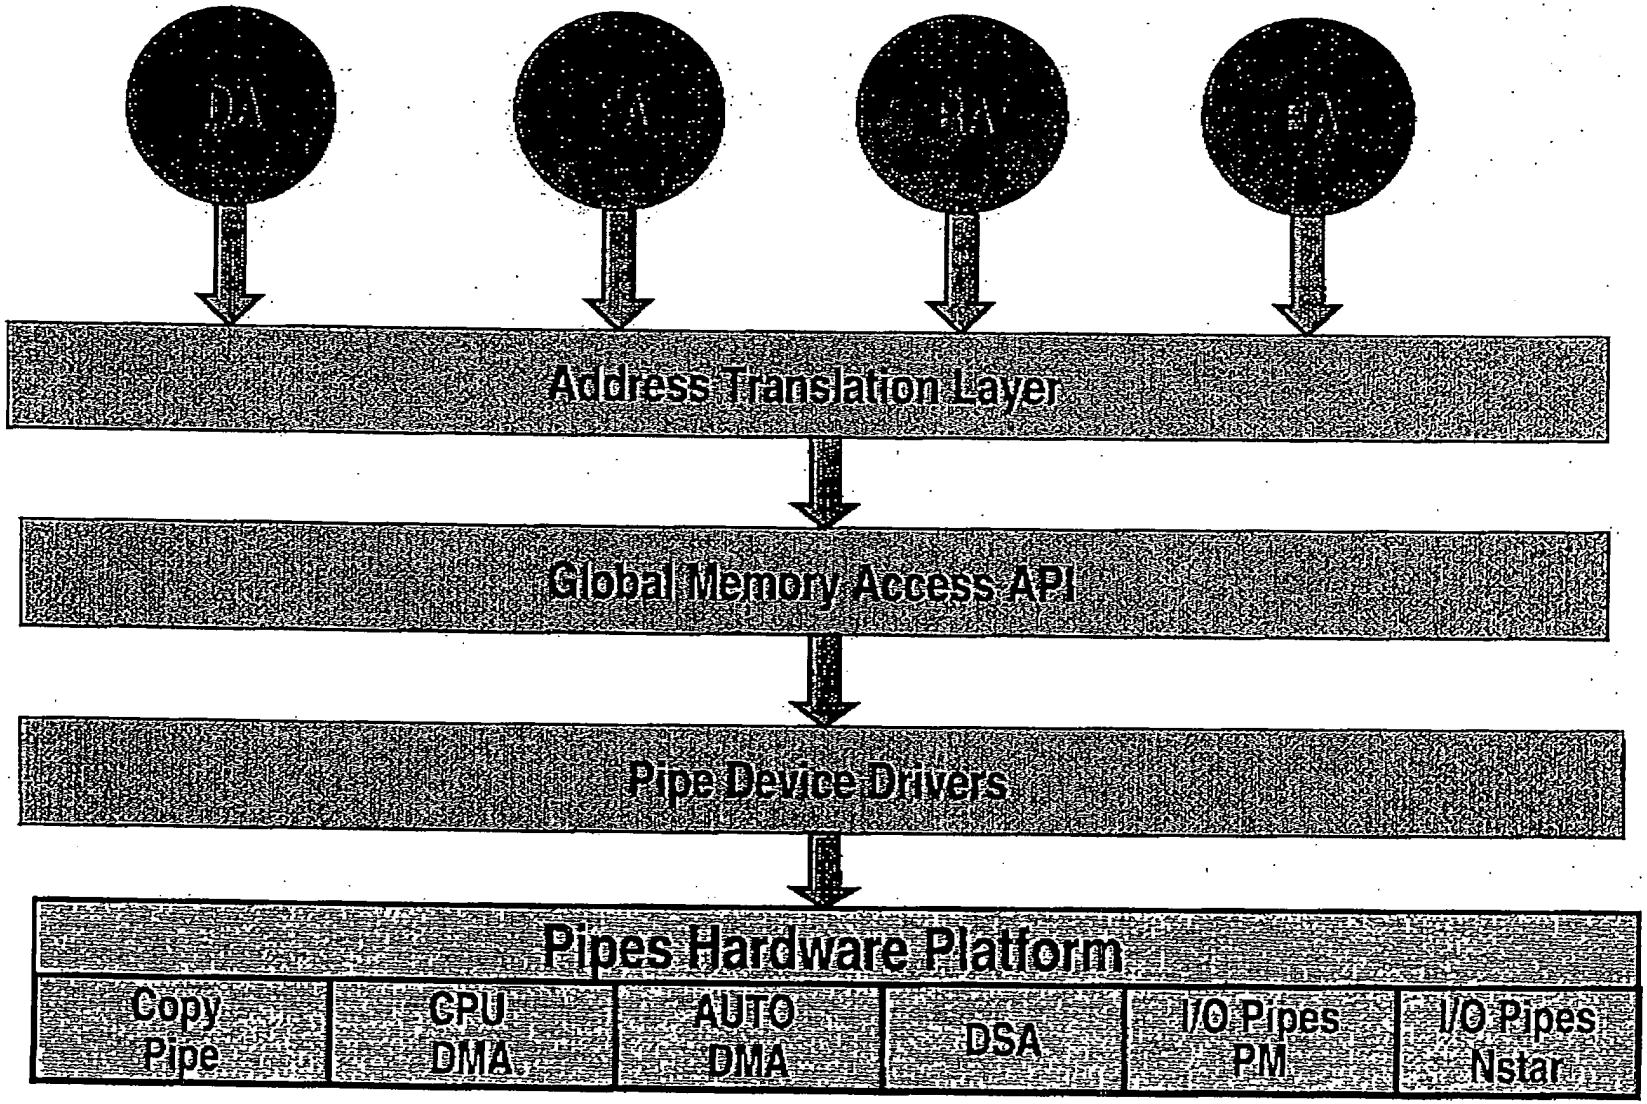

EMC²
where information lives

Memory Access Functional Layering

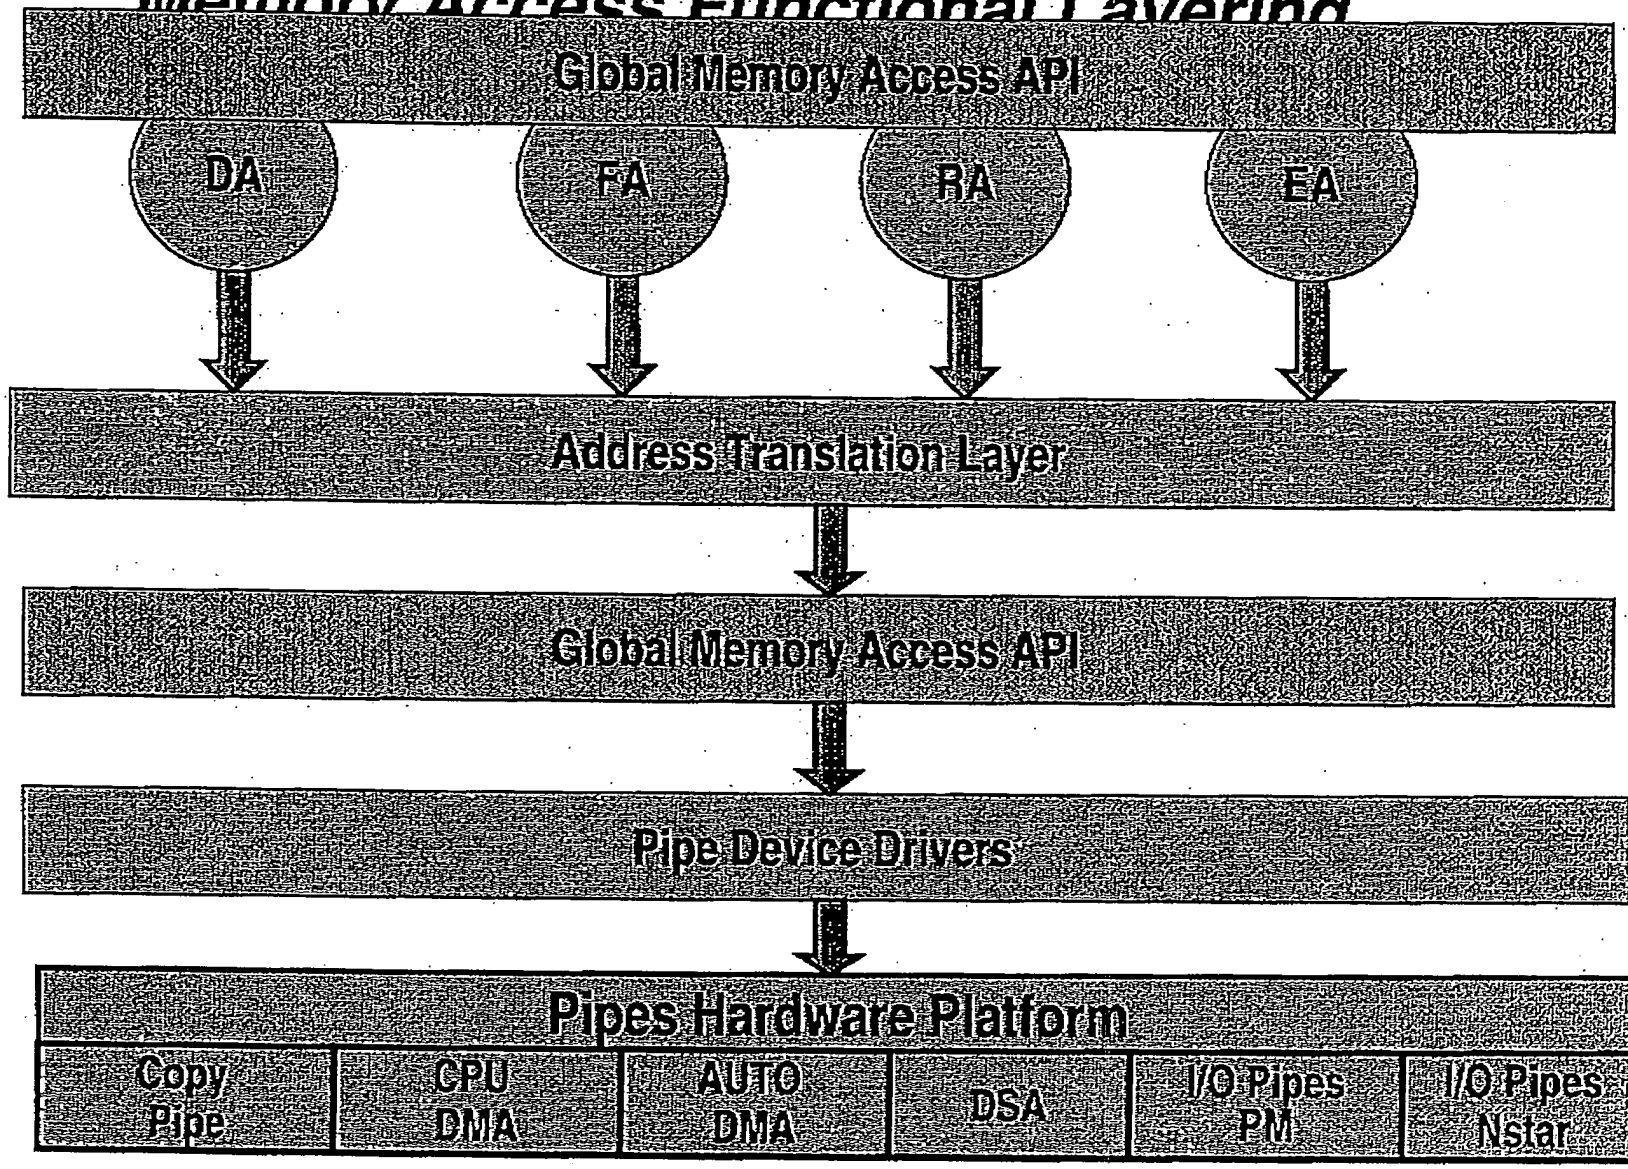

EMC²
where information lives

Global Memory Access API

- Manages all calls and return values from driver
- All error checking and error parsing
- Retries and basic error recovery
- Error & event tracing

All knowledge of mirrored memory lives at this layer

EMC²
where information lives

GMA Layer

- The Global Memory Access (GMA) layer will be responsible for maintaining the mirrored memory states.
- The GMA layer will tell the driver to write the data to both location (Normal) or to a single location (Degraded).
- The GMA layer will need to handle all issues with the atomic (MCMS) operations for mirrored memory when in both normal & degraded mode.
- The GMA layer is where the read load balancing algorithm lives.

Normal & Degraded State

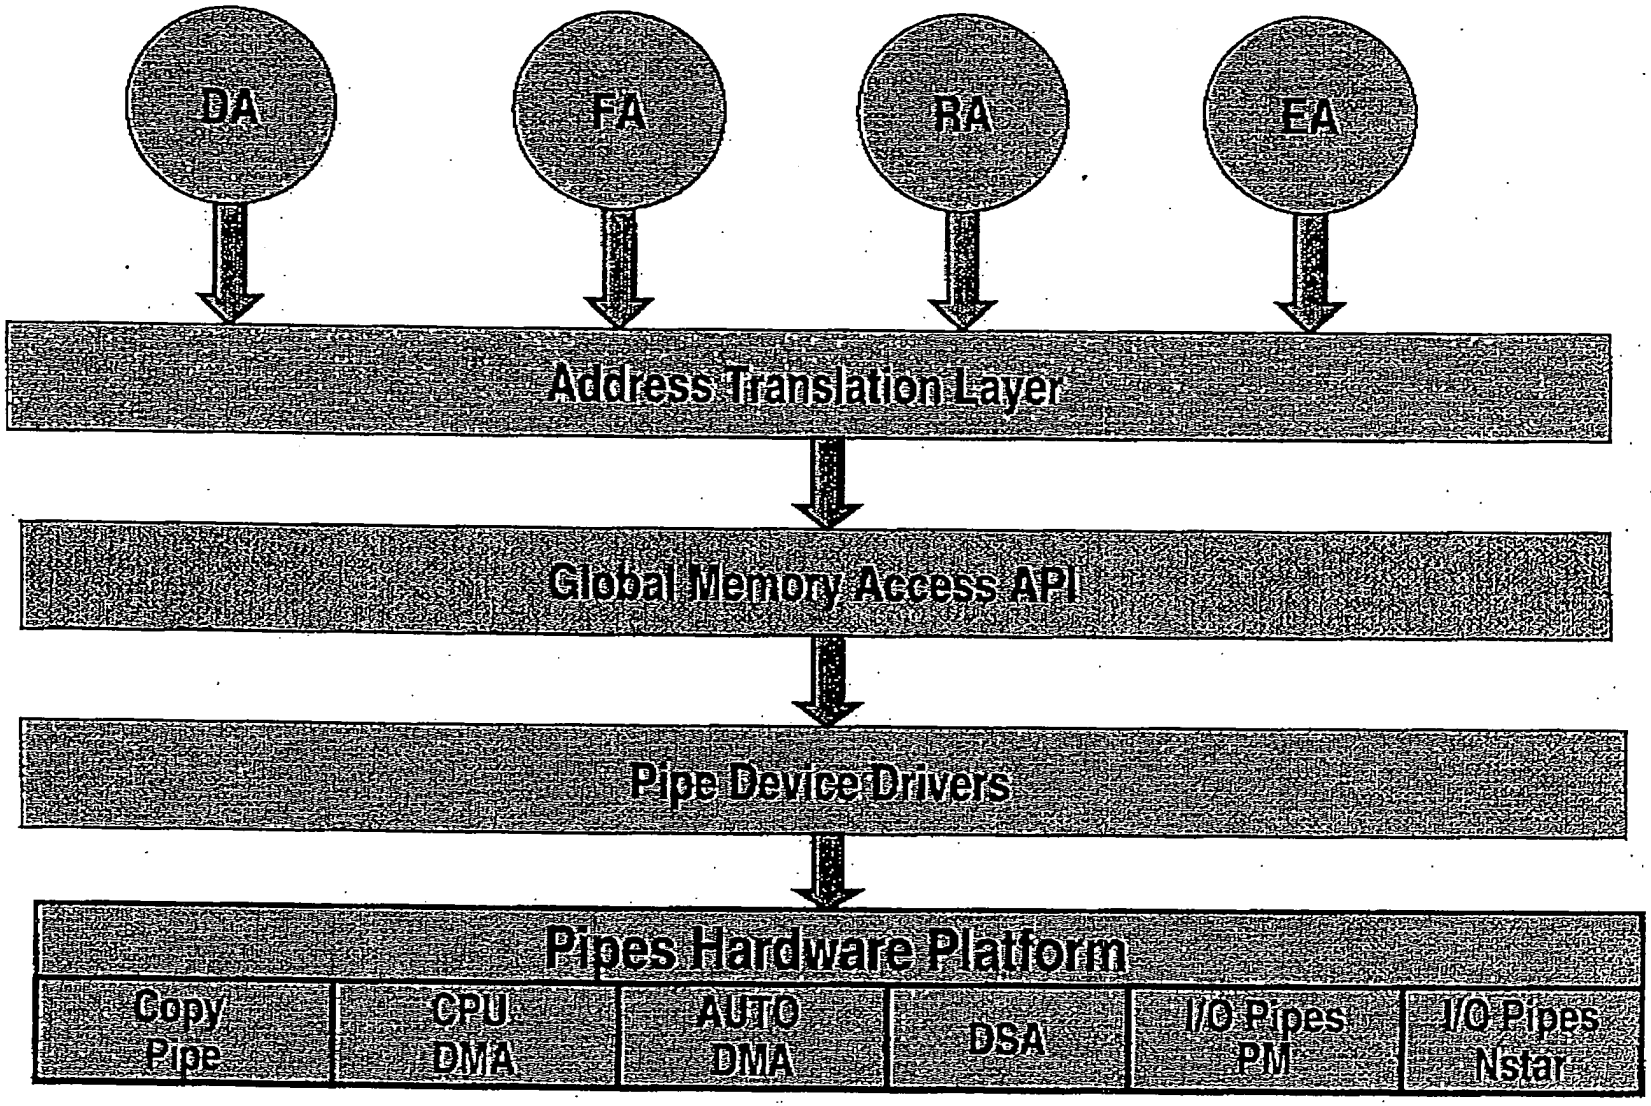

EMC²
where information lives

6 Different Mirrored Memory System States

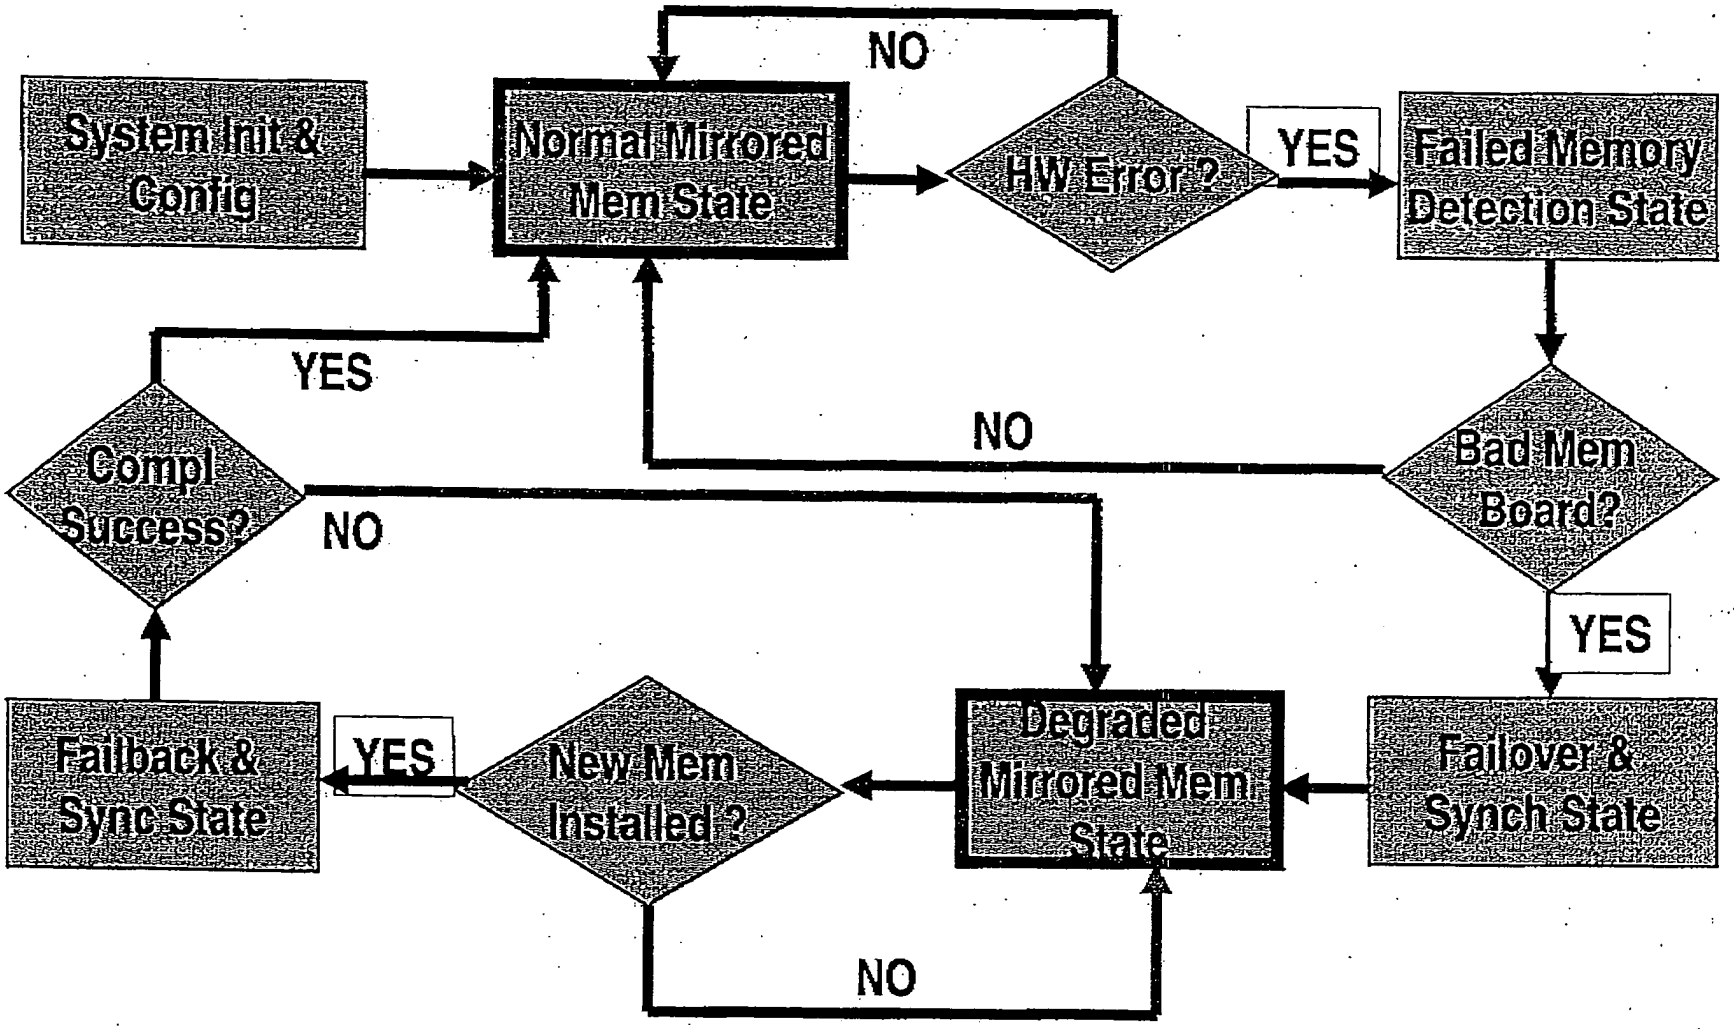

EMC²
where information lives

Failed Memory Detection State

- Main purpose of this phase is to diagnose a HW problem and decide if a memory board needs to be failed.
- The memory failure detection is encapsulated in the GMA layer.
- "Bad memory" detection is dependant on diagnosing where the problem resides.
- Some errors are obvious, but there are a subset of errors that it's difficult to determine if it's a memory problem.

EMC²
where information lives

Global Memory Access API

FAILED MEMORY DETECTION

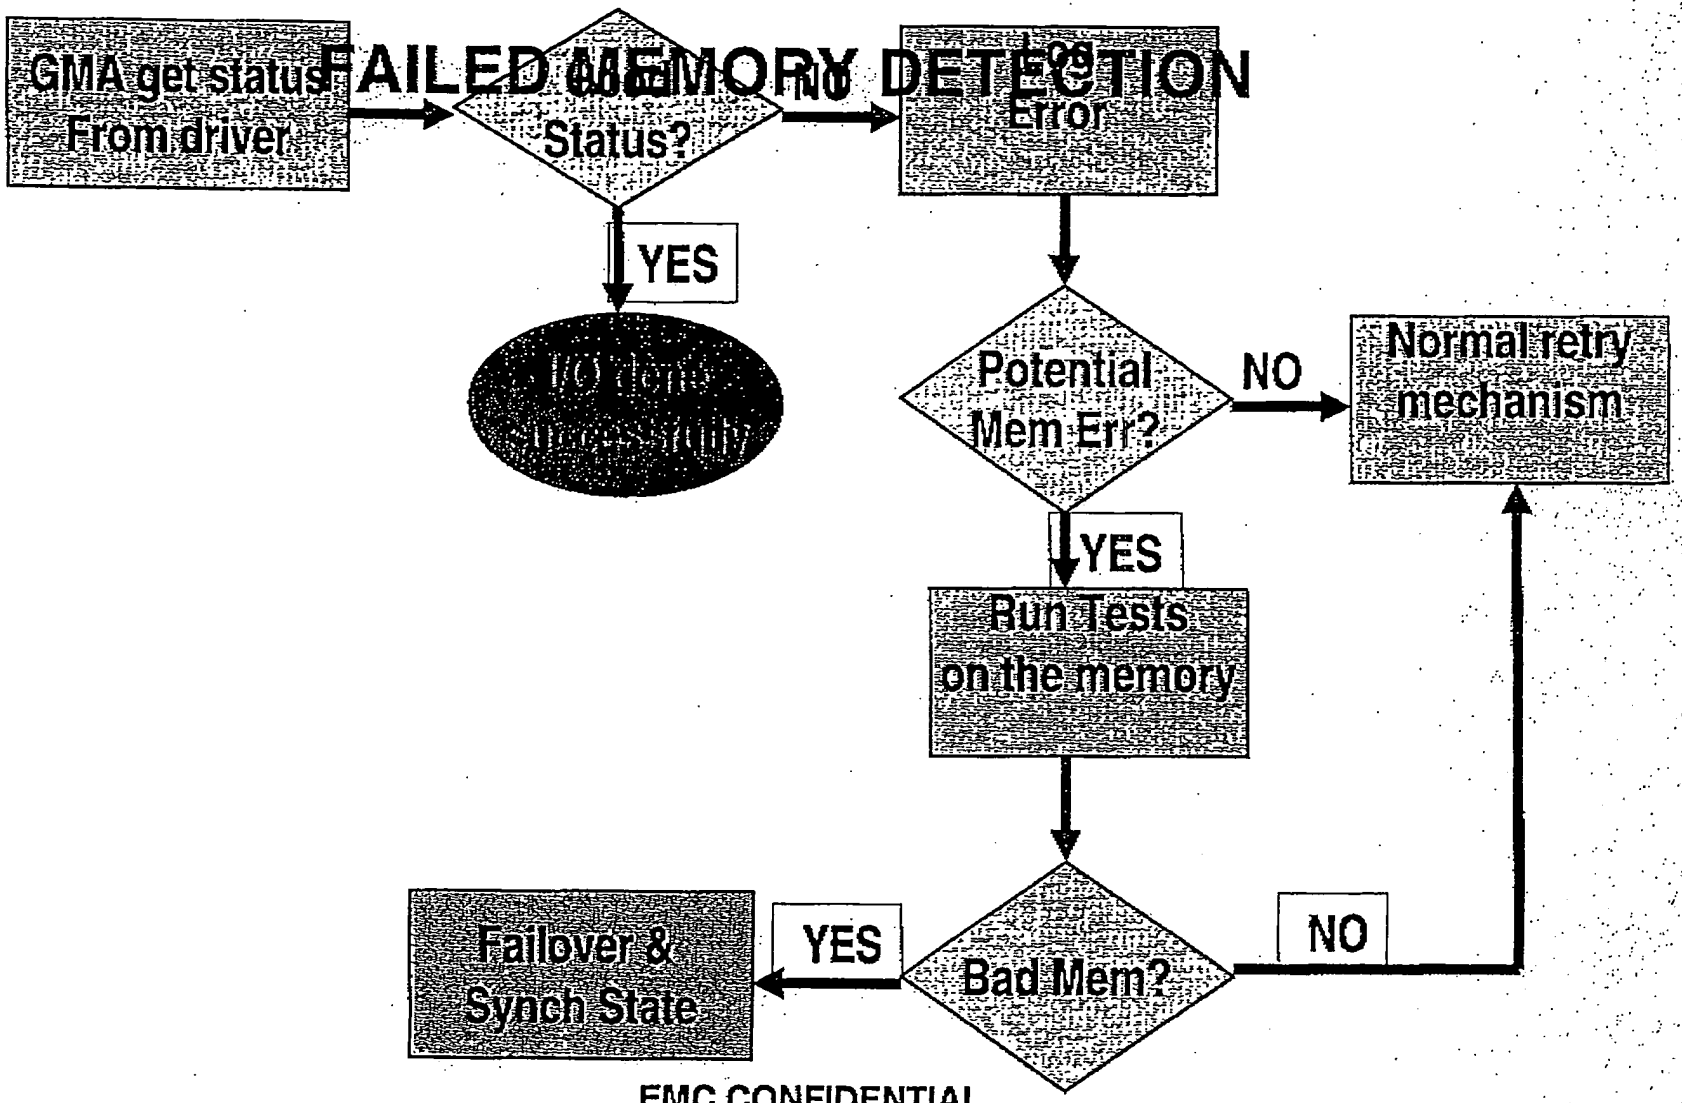

Failover & Synchronization State

- Software will use a level 6 interrupt (FMI) to synchronize failover between all directors.
- Software will use the message bus to communicate failover information between directors.
- There will be a separate "fail-over" thread on each director that will be invoked to perform the entire failover procedure.
- All new I/Os in the system will be on hold while the failover is taking place.

Failover & Synchronization State

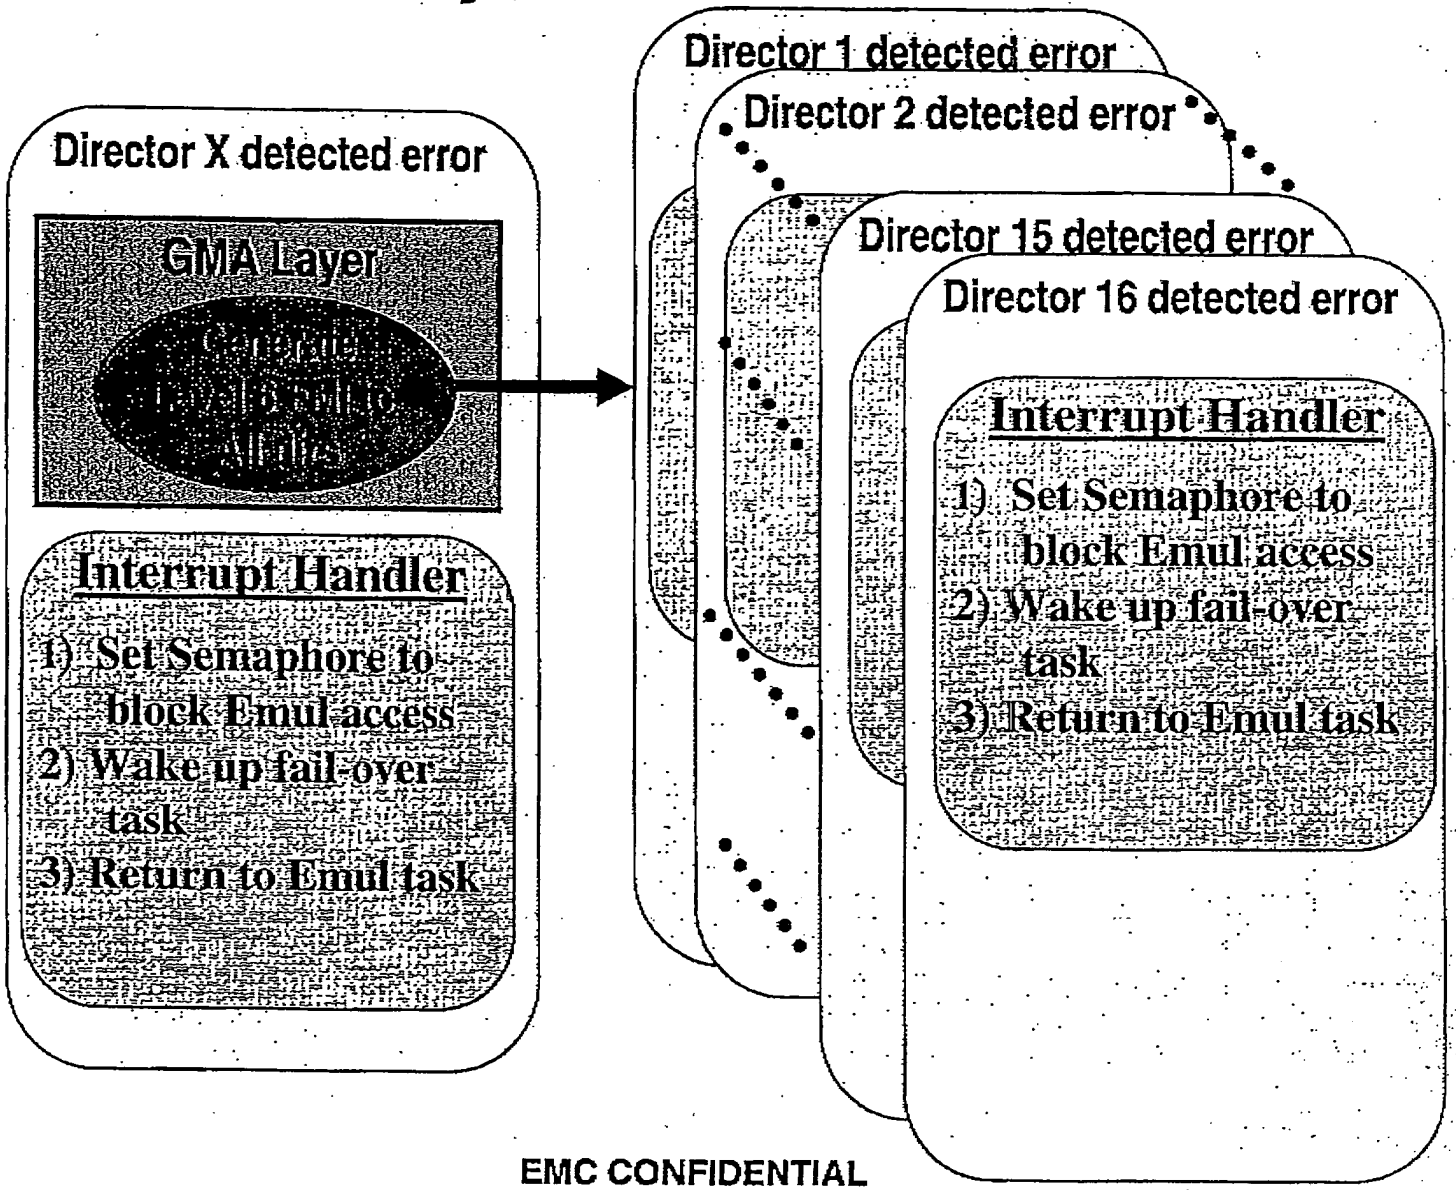

Failover & Synchronization State

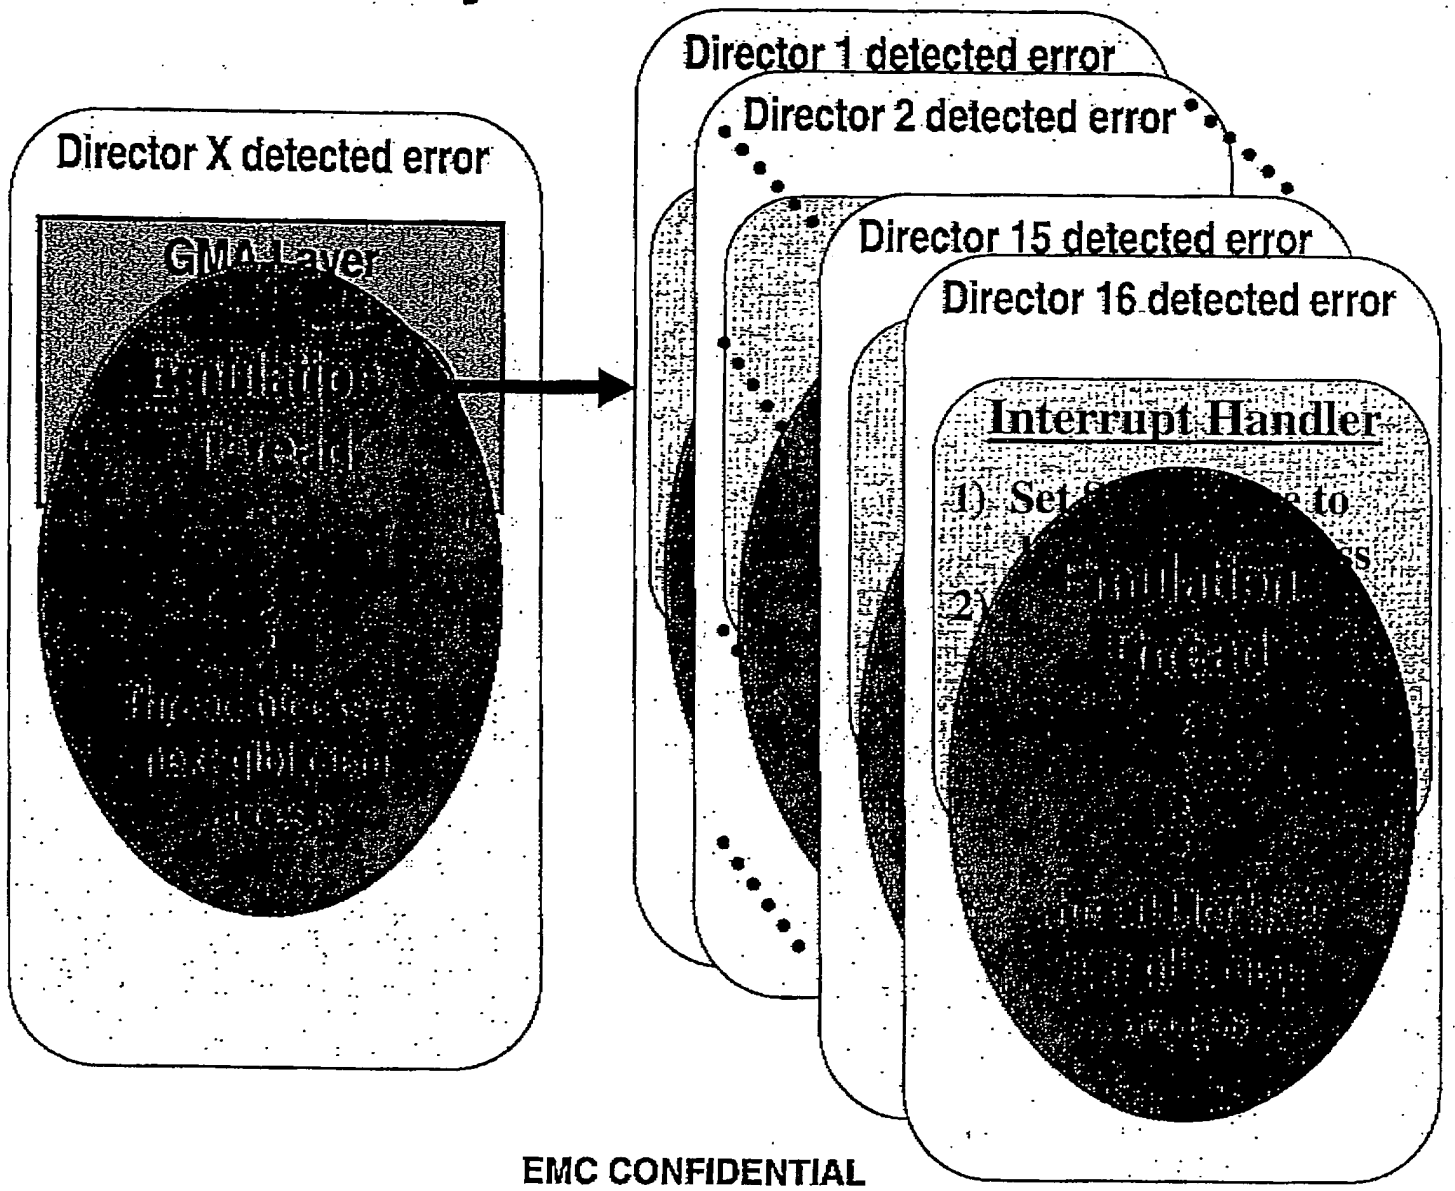

Failover & Synchronization State

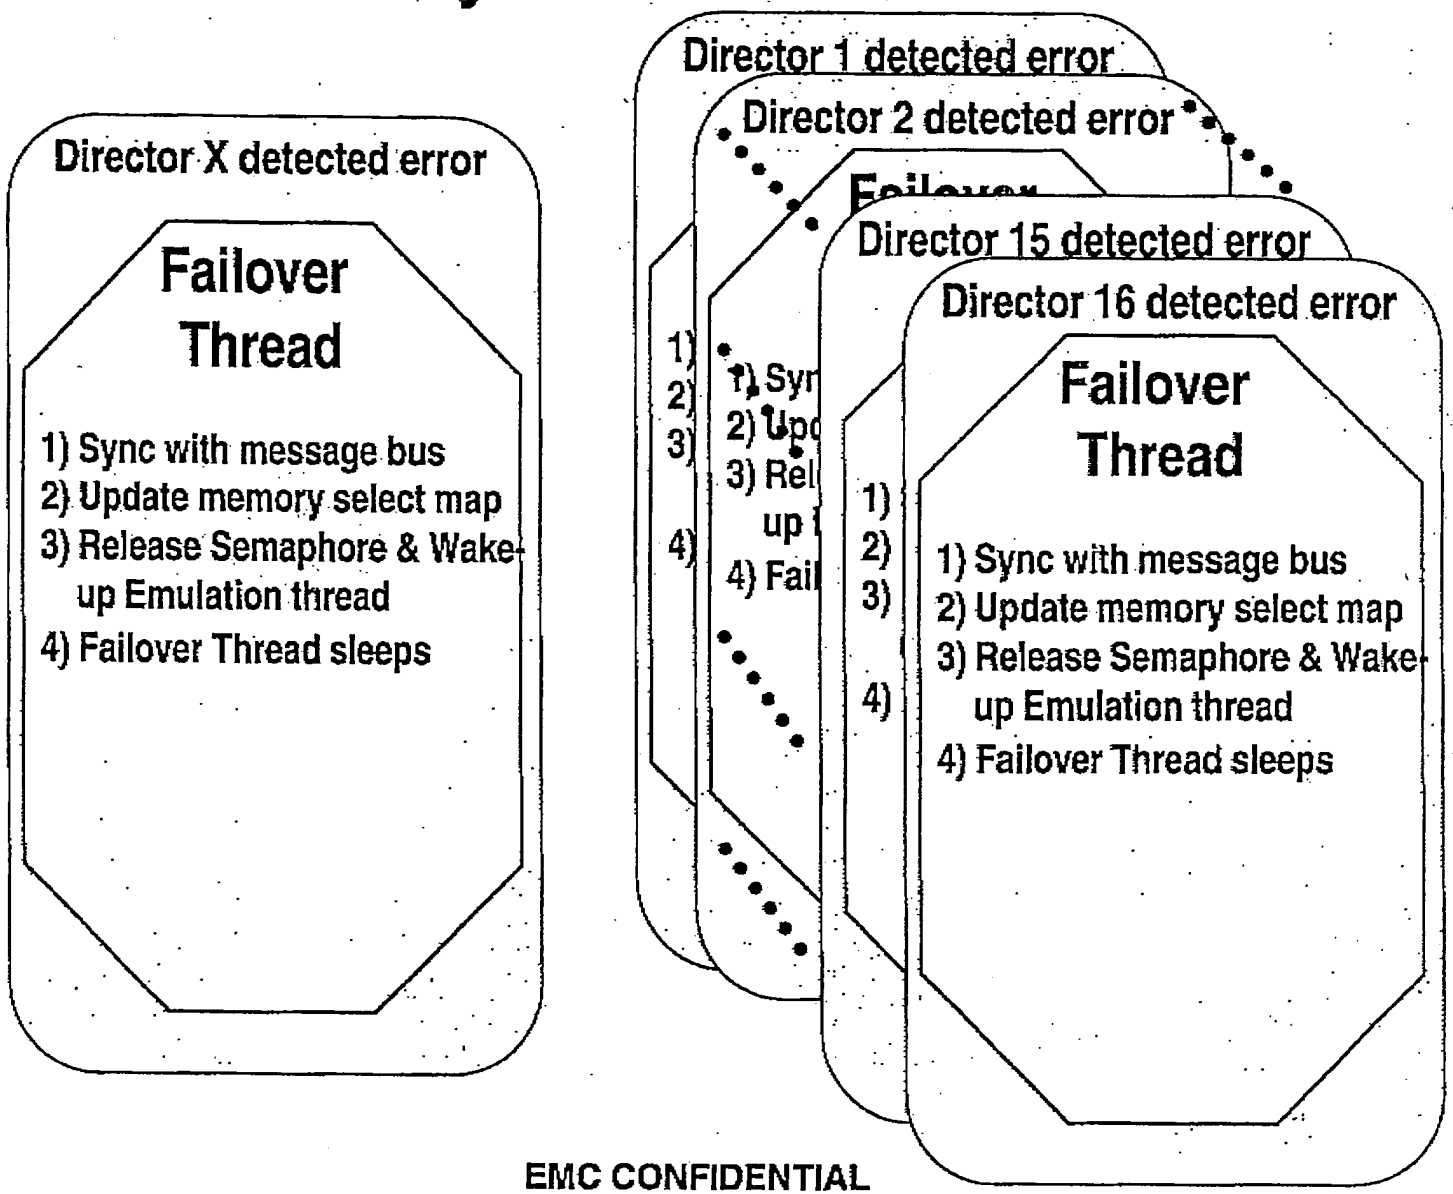

EMC²
where information lives

Failover & Synchronization State

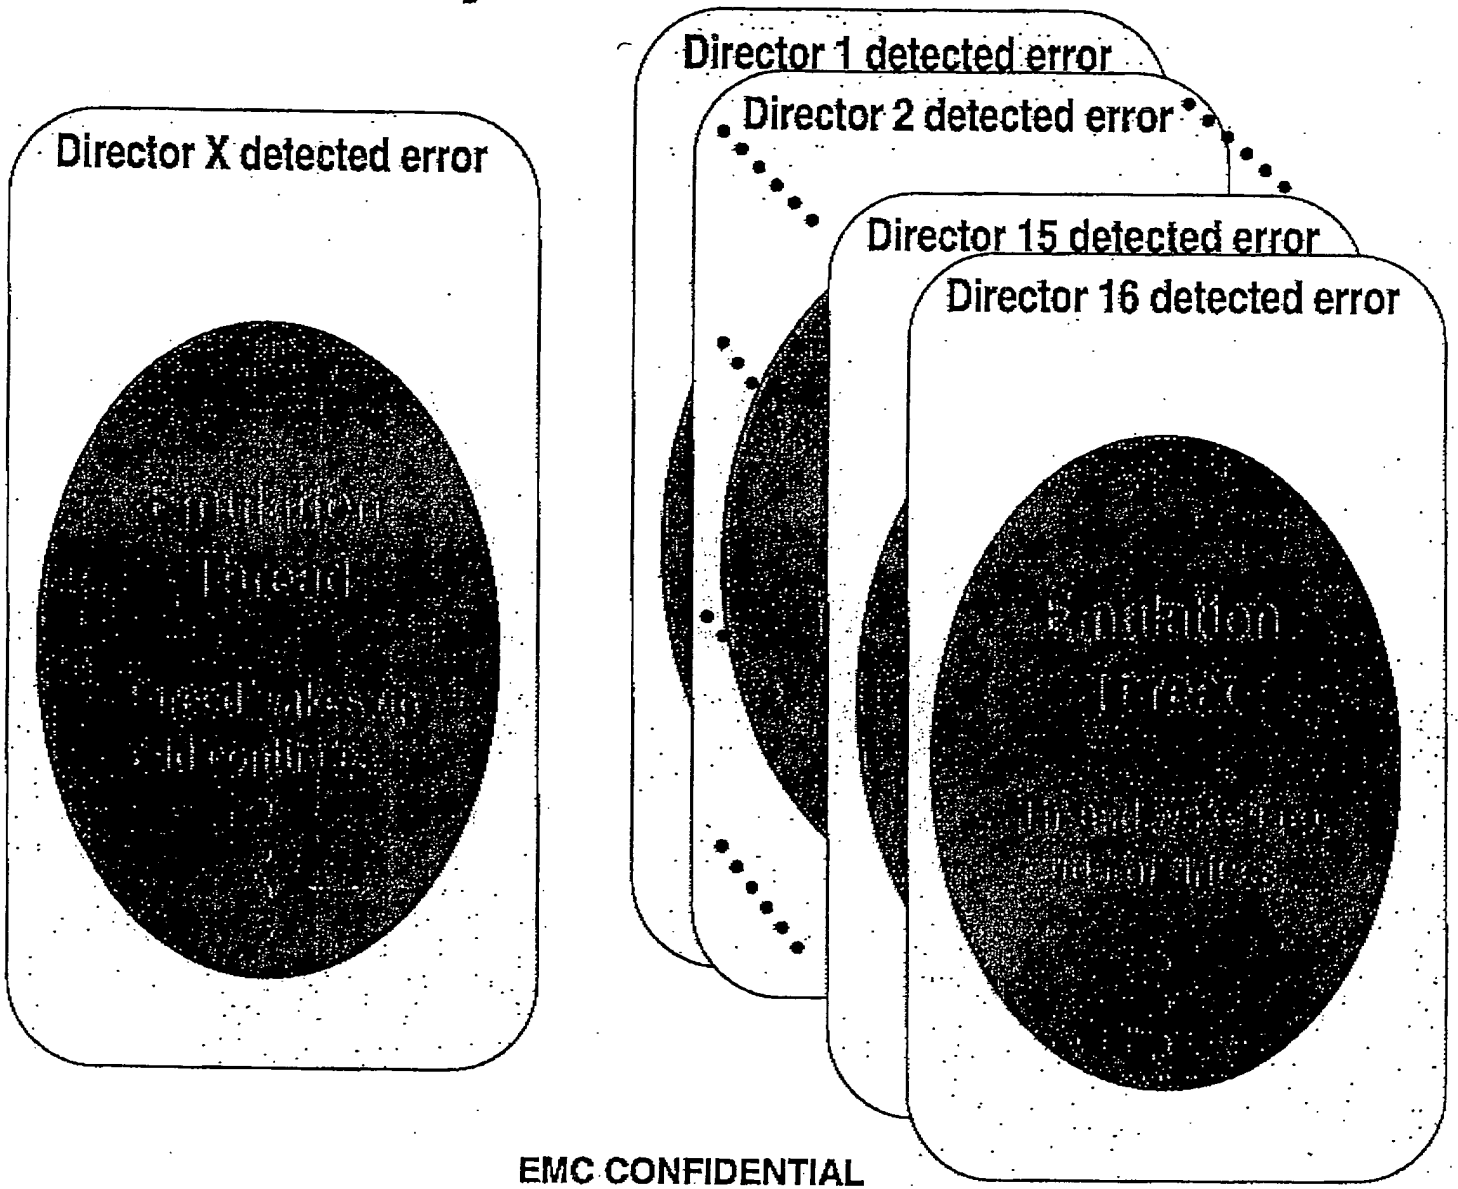

EMC²
where information lives

6 Different Mirrored Memory System States

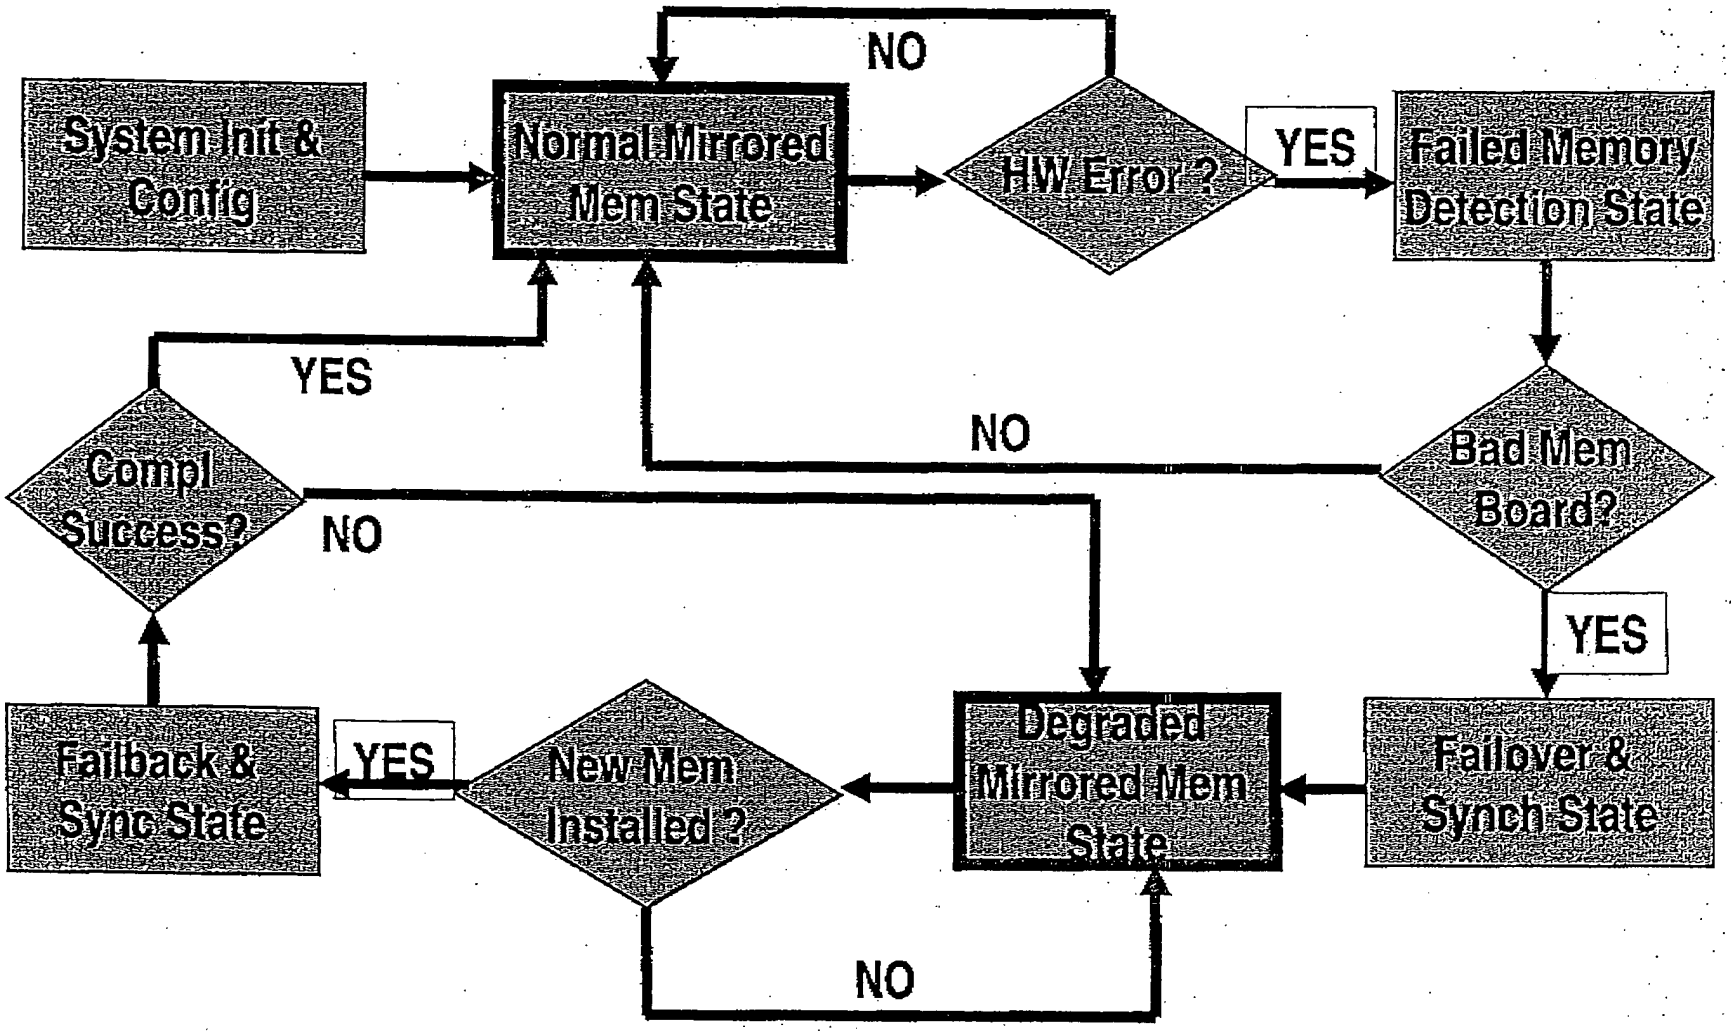

EMC²
where information lives

Failback & Synchronize State

- The procedure to failback to a new memory board is split into 3 separate phases
- All 3 phases can be done without any user or Symmwin intervention
- All directors need to coordinate on each of these phases.
- This procedure should have minimum to no impact on host I/O.
- The procedure should be simple and easy to use.

3 Phases of Failback Procedure

New Board

Memory Copy Phase

1. Wait for all active I/Os to complete (don't allow any new I/Os to begin)
2. Synchronize all directors
3. Each director copies a unit of memory (e.g. 1 Gig)
4. Synchronize all directors
5. Each director updates memory selection map
6. Release control back to emulation code for I/O execution
7. Let I/Os execute for a set time interval (e.g. 4 seconds) then repeat from step 1

EMC²
where information lives

6 Different Mirrored Memory System States

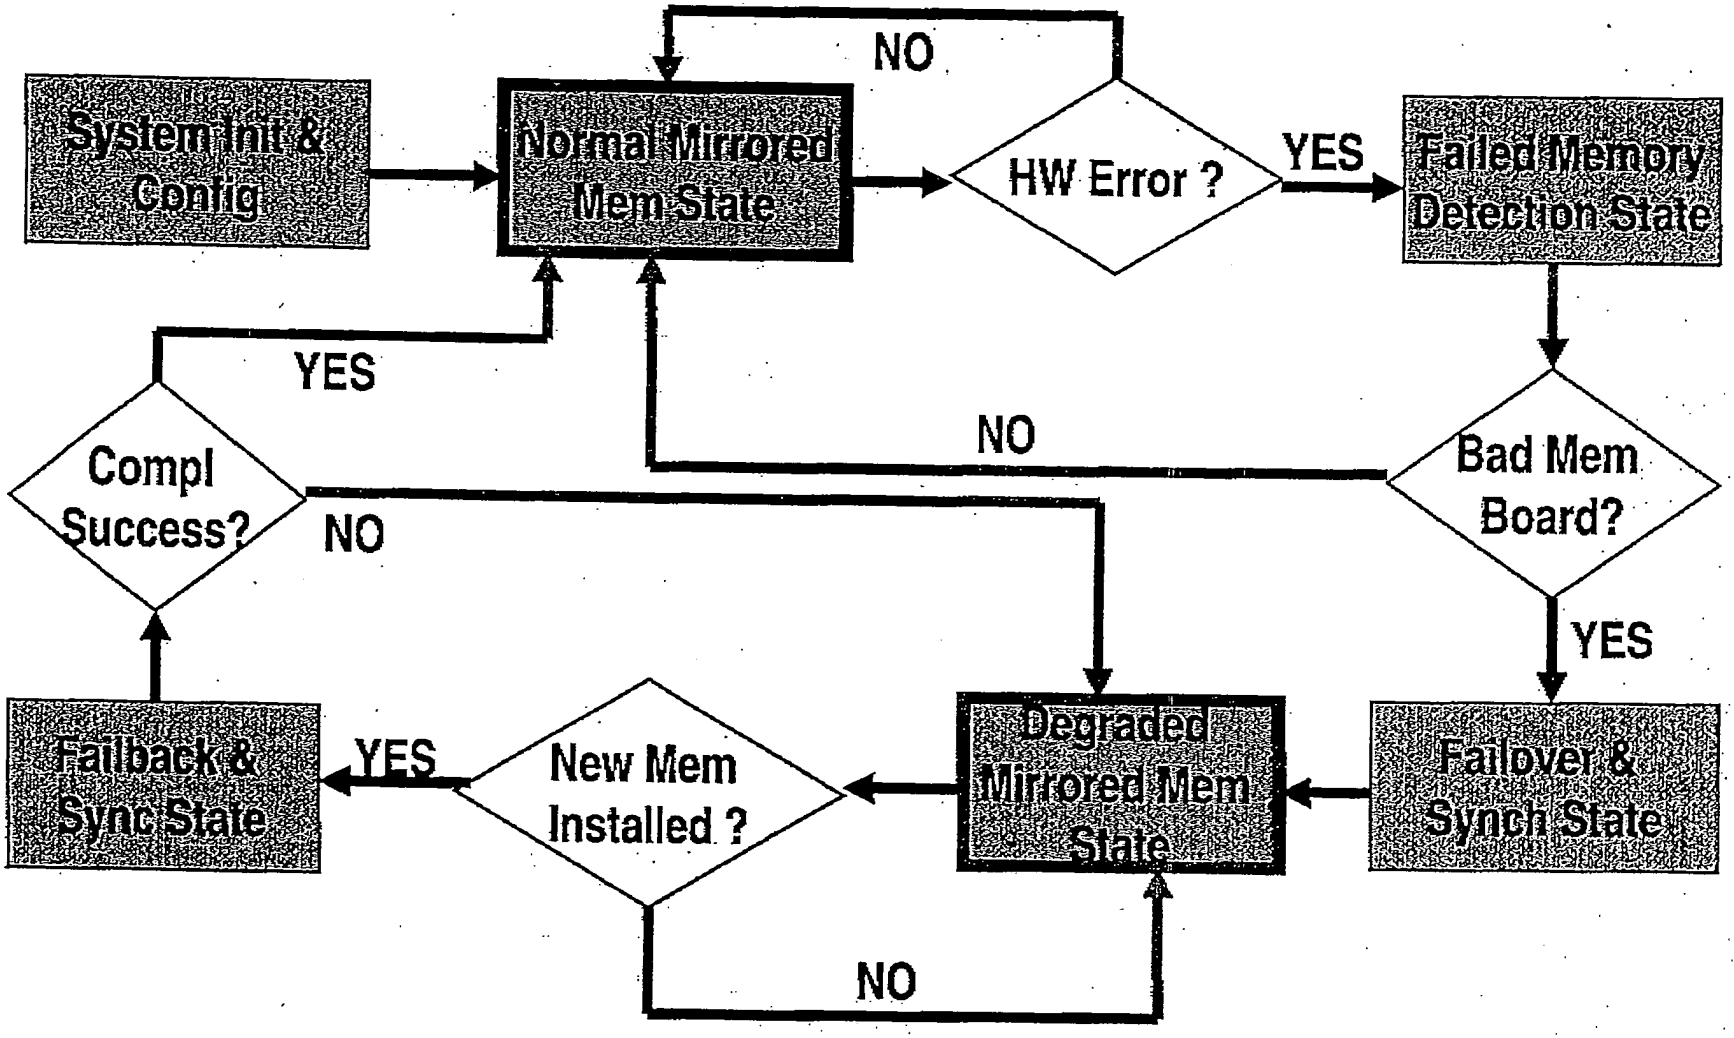

EMC²
where information lives

Mirror Memory Project Status

- Completed the Software Requirements Document (SRS)
- Started the Functional Spec
- Deep into design discussions
- Collecting design concepts and flowcharts for the High Level Design spec (HLDS)
- Developed a symm-6 prototype to demonstrate these concepts

EMC²

where information lives

EXHIBIT B

Subject: RE: Mirrored Memory External Review
Date:
From: araskiewicz, susan <araskiewicz_susan@emc.com>
To: cartmell, jerry <cartmell_jerry@emc.com>, decrescenzo, bob <decrescenzo_bob@emc.com>, rosner, tim <rosner_tim@emc.com>
CC: mcclure, steven <smcclure@emc.com>

Great I will set it up. Who should be included in the meeting?

Thanks
Susan

> -----Original Message-----
> From: cartmell, jerry
> Sent:
> To: decrescenzo, bob; araskiewicz, susan; rosner, tim
> Cc: mcclure, steven
> Subject: RE: Mirrored Memory External Review
>
> We can do a external review on the FS on (cs.)
>
> Jerry
>
> -----Original Message-----
> From: decrescenzo, bob
> Sent:
> To: araskiewicz, susan; rosner, tim; cartmell, jerry
> Cc: mcclure, steven
> Subject: RE: Mirrored Memory External Review
>
> I believe the fuctional spec is in good shape, but this is up to
> Jerry. He should be able to give you the state of the functional spec and
> also when he would be ready for an external review.
>
> --Bob
>
> -----Original Message-----
> From: araskiewicz, susan
> Sent:
> To: decrescenzo, bob; rosner, tim; cartmell, jerry
> Subject: RE: Mirrored Memory External Review
>
> OK, so why don't I plan on scheduling an External Functional
> Review.
>

> Bob does the document need to be revised or is it ready to
> be sent out for review before the external meeting. I can target setting
> the meeting up for next week if that works.
>
> Tim who needs to be on the external distribution list?
>
> Thanks
> Susan
>
> -----Original Message-----
> From: decrescenzo, bob
> Sent:
> To: araskiewicz, susan; rosner, tim; cartmell,
> jerry
> Subject: RE: Mirrored Memory External Review
>
> Yes. We only had an internal review. There has
> been nothing done externally.
>
> -----Original Message-----
> From: araskiewicz, susan
> Sent:
> To: decrescenzo, bob; rosner, tim; cartmell,
> jerry
> Subject: RE: Mirrored Memory External Review
>
> Bob, I was under the impression you already had the
> functional review. Was it just an internal review?
>
> Thanks
> Susan
>
> -----Original Message-----
> From: decrescenzo, bob
> Sent:
> To: rosner, tim; araskiewicz, susan; cartmell,
> jerry
> Subject: RE: Mirrored Memory External Review
>
> I think they will get this from a functional review.
> I don't think anyone can get anything out of design review without having
> a functional review first. The functional review would explain "what" we
> are going to be doing. The design review explains "how" we did it
> internally in the code. If they don't thoroughly know and understand the
> "what" then the "how" is going to be meaningless.
>
> --Bob
>
> -----Original Message-----
> From: rosner, tim
> Sent:
> To: decrescenzo, bob; araskiewicz, susan;
> cartmell, jerry
> Subject: RE: Mirrored Memory External Review
>
> Bob,
>
> I understand your point, however if you don't mind I
> would like to continue with the external review. Trust me I will not
> allow any external group influence our design.

>
> The reason why I want external involvement is
> because I want to have them understand the extensive work your group is
> going threw and that it's not just a simple change.
>
> TR
>
> -----Original Message-----
> From: decrescenzo, bob
> Sent:
> To: araskiewicz, susan; cartmell, jerry
> Cc: rosner, tim
> Subject: RE: Mirrored Memory External Review
>
> Susan,
>
> I don't know if we want to have an external
> design review. The only information that external groups should care
> about is in the functional spec. We may want to have an external
> functional review. I don't think product management, marketing, etc.
> really care if we are using separate internal threads to do fail-over or
> if we do this fail-over inside existing threads. This is the details in
> the design spec. They really only care about how the feature going to
> "function". If design questions come up we can explain how we are going
> to implement something, but I don't think this should be the focus to
> external groups.
>
> We don't want to give external groups the ability to
> direct/change our designs. We want their feedback on how the feature is
> going to function to make sure it's in alignment with what they expected.
> I don't think this applies to our own internal designs.
>
> Tim, what do you think?
>
>
>
> -----Original Message-----
> From: araskiewicz, susan
> Sent:
> To: cartmell, jerry
> Cc: decrescenzo, bob
> Subject: Mirrored Memory External Review
>
> Hi Jerry, Congratulation on a very successful
> Internal Design Review for Mirrored Memory.
>
> Below are my meeting minutes from the Mirrored
> Memory Internal Design Review.
>
>
>
> Action Items:
>
> Need to come to an agreement with Vault. (Jerry)
>
> Wayne Sylvia will send question to Jerry. (Wayne)
>
> Discuss Special Command with Haim (Jerry)
>
>

> Discussions:
>
> You still want SymWin Scripts
>
> Scott Gordon will write a utility (Syscall) to light
> the light (LED)
>
>
> When would you like to schedule and External Review.
> I want to give you enough time to complete the action items.
>
> Thanks
> Susan